



US009448847B2

(12) **United States Patent**  
**Sandstrom**

(10) **Patent No.:** **US 9,448,847 B2**  
(45) **Date of Patent:** **Sep. 20, 2016**

(54) **CONCURRENT PROGRAM EXECUTION  
OPTIMIZATION**

G06F 9/3851; G06F 9/3887; G06F 9/3891;  
H04L 67/1038

See application file for complete search history.

(71) Applicant: **Mark Henrik Sandstrom**, Jersey City,  
NJ (US)

(56) **References Cited**

(72) Inventor: **Mark Henrik Sandstrom**, Jersey City,  
NJ (US)

U.S. PATENT DOCUMENTS

(73) Assignee: **THROUGHPUTER, INC.**, Jersey City,  
NJ (US)

(\*) Notice: Subject to any disclaimer, the term of this  
patent is extended or adjusted under 35  
U.S.C. 154(b) by 129 days.

(21) Appl. No.: **14/318,512**

(22) Filed: **Jun. 27, 2014**

(65) **Prior Publication Data**

US 2015/0058857 A1 Feb. 26, 2015

**Related U.S. Application Data**

(60) Provisional application No. 61/934,747, filed on Feb.  
1, 2014, provisional application No. 61/869,646, filed  
on Aug. 23, 2013.

(51) **Int. Cl.**  
**H04J 3/04** (2006.01)  
**G06F 9/50** (2006.01)

(52) **U.S. Cl.**  
CPC ..... **G06F 9/5038** (2013.01); **G06F 2209/5021**  
(2013.01)

(58) **Field of Classification Search**  
CPC ..... G06F 9/5038; G06F 2209/5021;  
G06F 9/505; G06F 2206/1012; G06F 9/3012;

7,447,873 B1 \* 11/2008 Nordquist ..... G06F 9/30043  
712/2  
8,539,207 B1 \* 9/2013 LeGrand ..... G06F 15/167  
712/225  
2011/0321057 A1 \* 12/2011 Mejdrich ..... G06F 9/505  
718/105  
2012/0324458 A1 \* 12/2012 Peterson ..... G06F 9/5038  
718/102  
2013/0222402 A1 \* 8/2013 Peterson ..... G06F 9/52  
345/520  
2014/0317378 A1 \* 10/2014 Lippett ..... G06F 1/3203  
712/15  
2015/0339798 A1 \* 11/2015 Peterson ..... G06F 9/52  
345/520  
2015/0378776 A1 \* 12/2015 Lippett ..... G06F 1/3203  
718/101

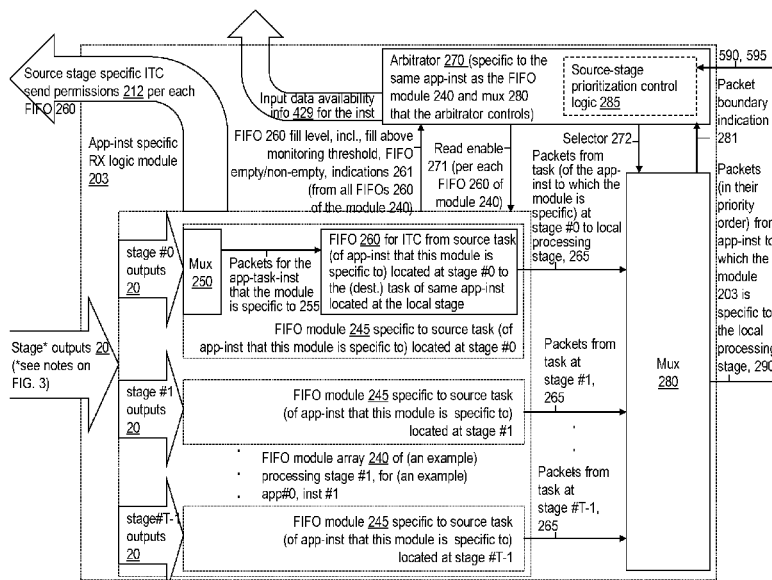
\* cited by examiner

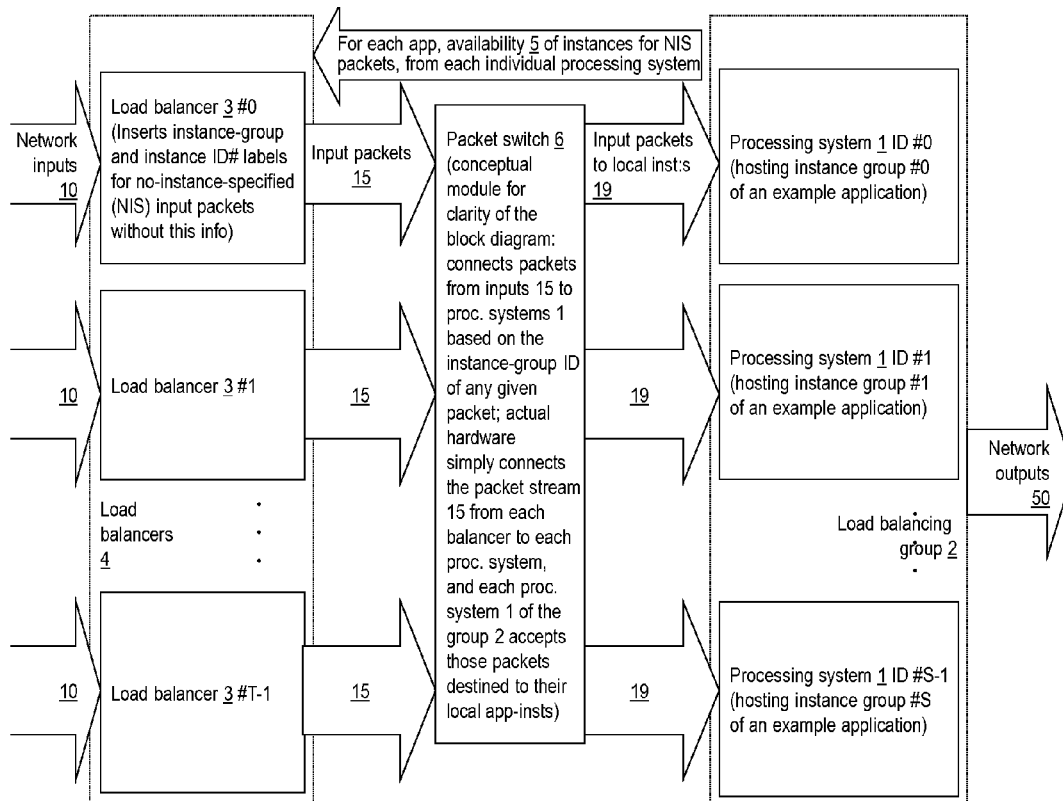
*Primary Examiner* — Brian O'Connor

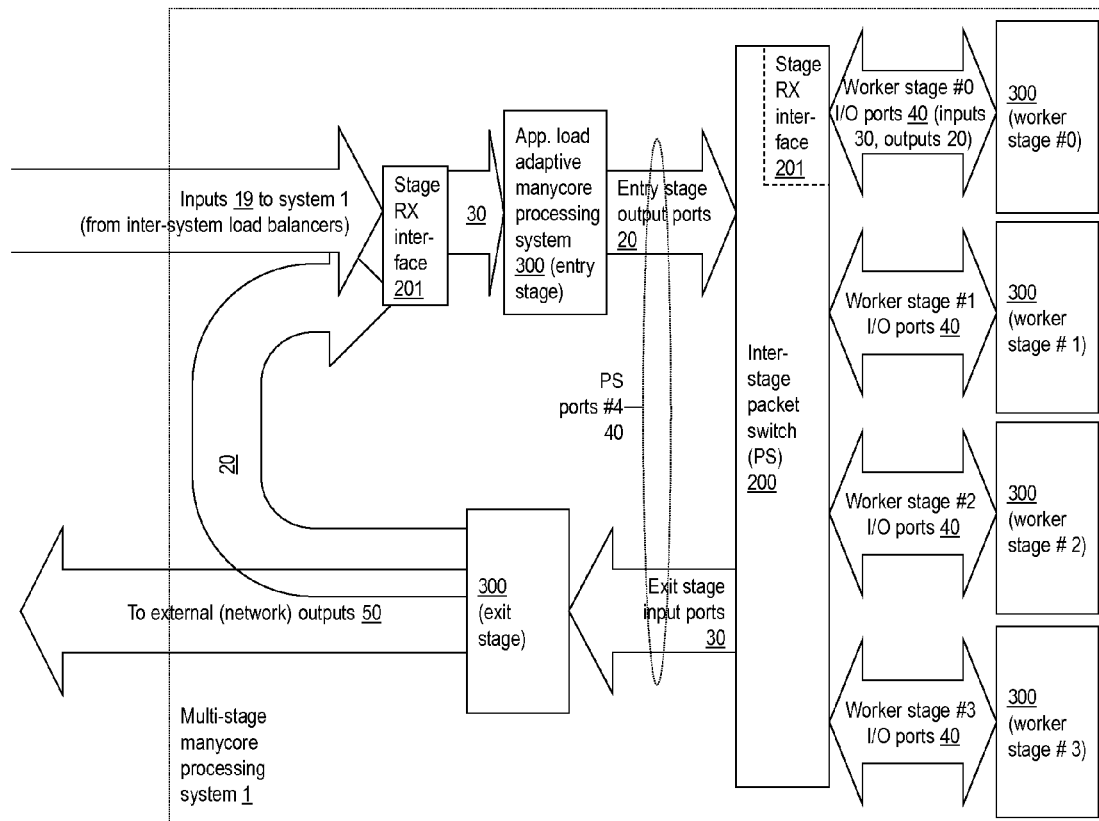
(57) **ABSTRACT**

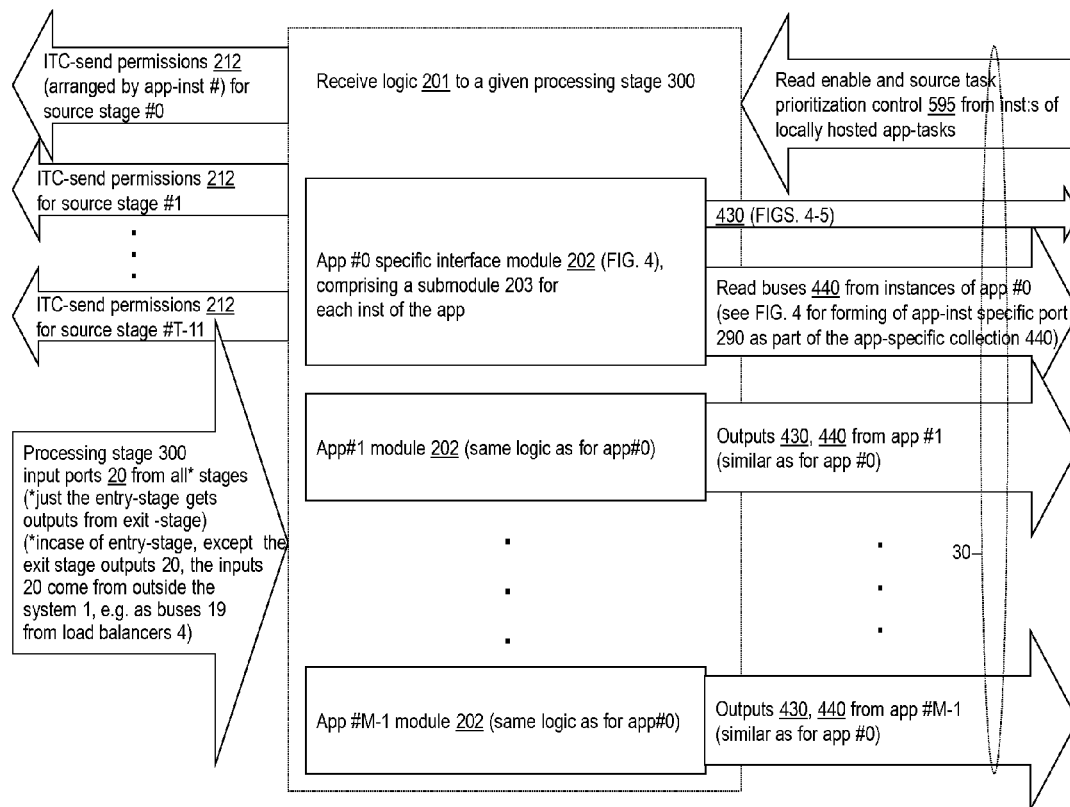
An architecture for a load-balanced groups of multi-stage  
manycore processors shared dynamically among a set of  
software applications, with capabilities for destination task  
defined intra-application prioritization of inter-task commu-  
nications (ITC), for architecture-based ITC performance  
isolation between the applications, as well as for prioritizing  
application task instances for execution on cores of many-  
core processors based at least in part on which of the task  
instances have available for them the input data, such as ITC  
data, that they need for executing.

**18 Claims, 7 Drawing Sheets**



**FIG. 1**

**FIG. 2**

**FIG. 3**

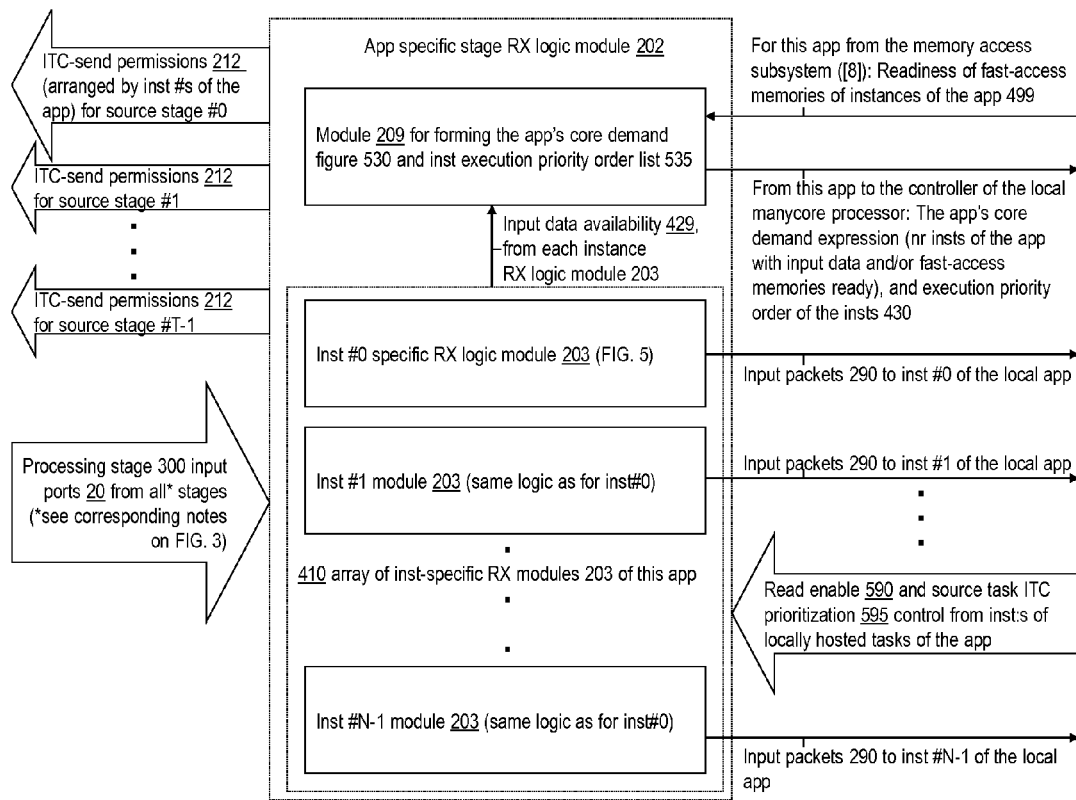


FIG. 4

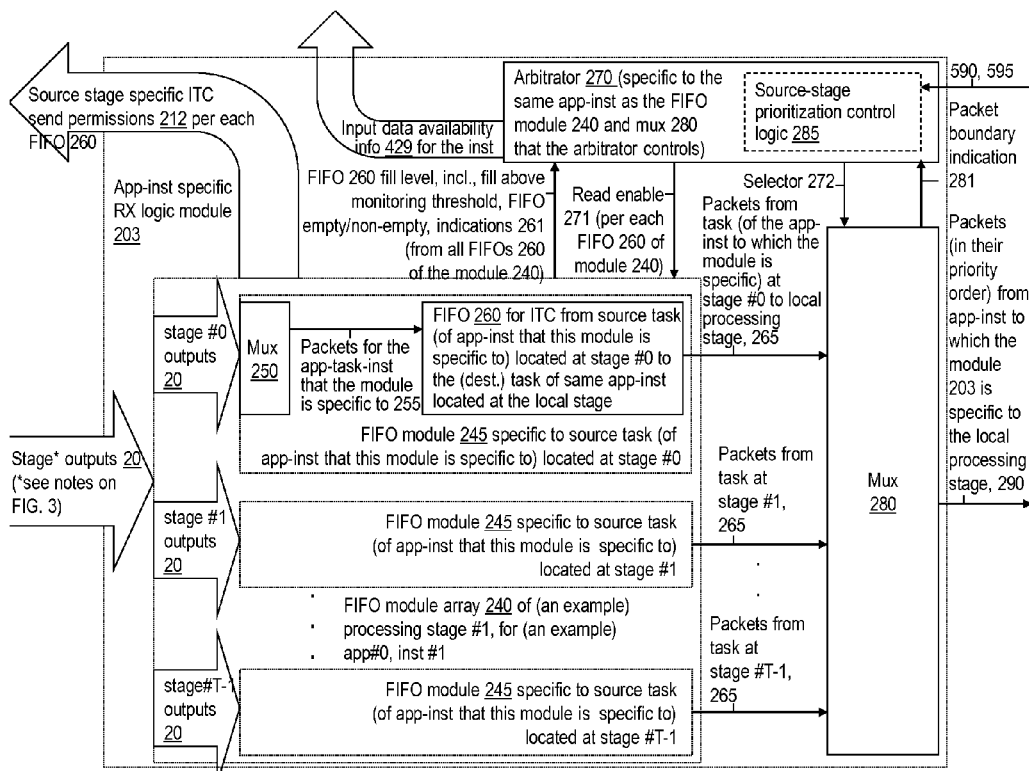


FIG. 5

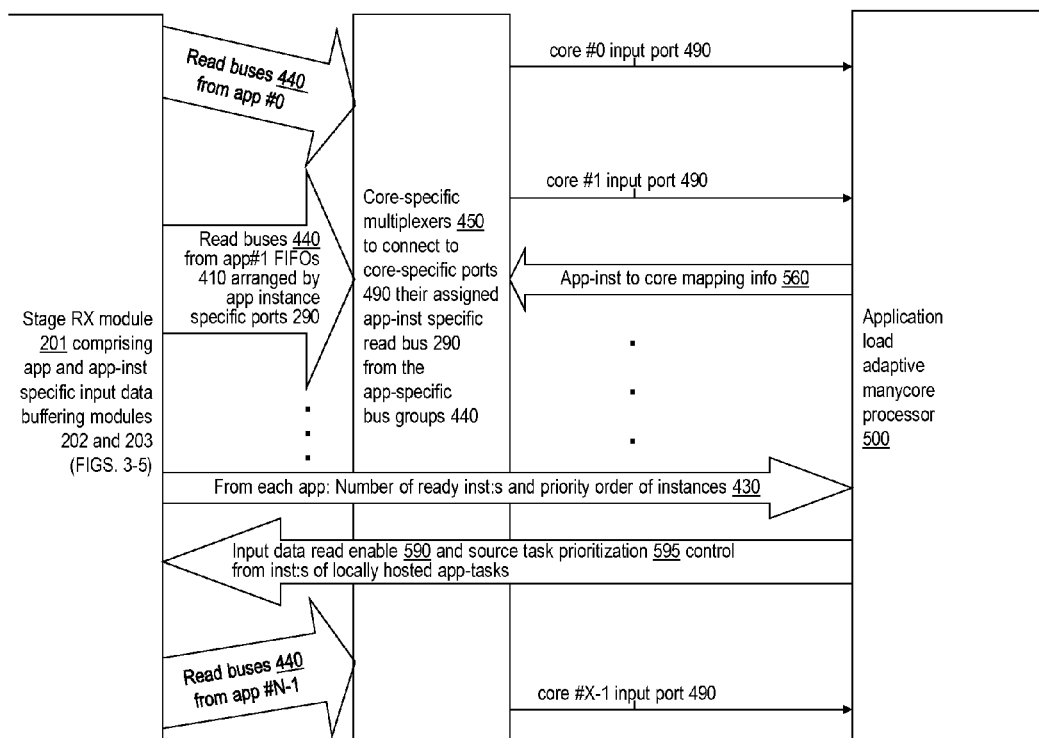


FIG. 6

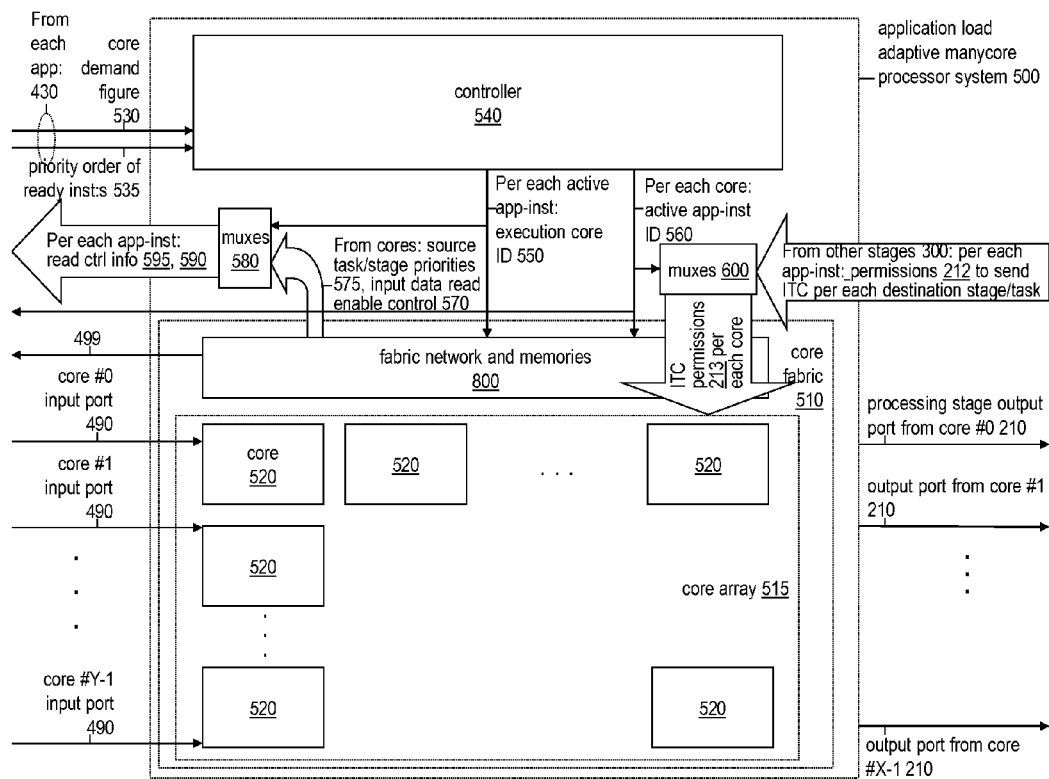


FIG. 7



1

## CONCURRENT PROGRAM EXECUTION OPTIMIZATION

### CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims the benefit of the following applications, each of which is incorporated by reference in its entirety:

- [1] U.S. Provisional Application No. 61934747, filed Feb. 1, 2014;
- [2] U.S. Provisional Application No. 61869646, filed Aug. 23, 2013; and
- [3] U.S. Utility application Ser. No. 13/959,596, filed Aug. 5, 2013.

This application is also related to the following co-pending or patented applications, each of which is incorporated by reference in its entirety:

- [4] U.S. Utility application Ser. No. 13/184,028, filed Jul. 15, 2011;
- [5] U.S. Utility application Ser. No. 13/270,194, filed Oct. 10, 2011;
- [6] U.S. Utility application Ser. No. 13/277,739, filed Nov. 21, 2011;
- [7] U.S. Utility application Ser. No. 13/297,455, filed Nov. 16, 2011;
- [8] U.S. Utility application Ser. No. 13/684,473, filed Nov. 23, 2012;
- [9] U.S. Utility application Ser. No. 13/717,649, filed Dec. 17, 2012;
- [10] U.S. Utility application Ser. No. 13/901,566, filed May 24, 2013; and
- [11] U.S. Utility application Ser. No. 13/906,159, filed May 30, 2013.

### BACKGROUND

#### 1. Technical Field

This invention pertains to the field of information processing, particularly to techniques for managing execution of multiple concurrent, multi-task software programs on parallel processing hardware.

#### 2. Descriptions of the Related Art

Conventional microprocessor and computer system architectures rely on system software for handling runtime matters relating to sharing processing resources among multiple application programs and their instances, tasks etc., as well as orchestrating the concurrent (parallel and/or pipelined) execution between and within the individual applications sharing the given set of processing resources. However, the system software consumes by itself ever increasing portions of the system processing capacity, as the number of applications, their instances and tasks and the pooled processing resources would grow, as well as the more frequently the optimizations of the dynamic resource management among the applications and their tasks would be needed to be performed, in response to variations in the applications' and their instances' and tasks' processing loads etc. variables of the processing environment. As such, the conventional approaches for supporting dynamic execution of concurrent programs on shared processing capacity pools will not scale well.

This presents significant challenges to the scalability of the networked utility ('cloud') computing model, in particular as there will be a continuously increasing need for greater degrees of concurrent processing also at intra-application levels, in order to enable increasing individual application

2

on-time processing throughput performance, without the automatic speed-up from processor clock rates being available due to the practical physical and economic constraints faced by the semiconductor etc. physical hardware implementation technologies.

To address the challenges per above, there is a need for inventions enabling scalable, multi-application dynamic concurrent execution on parallel processing systems, with high resource utilization efficiency, high application processing on-time throughput performance, as well built-in, architecture based security and reliability.

### SUMMARY

An aspect of the invention provides systems and methods for arranging secure and reliable, concurrent execution of a set of internally parallelized and pipelined software programs on a pool of processing resources shared dynamically among the programs, wherein the dynamic sharing of the resources is based at least in part on i) processing input data loads for instances and tasks of the programs and ii) contractual capacity entitlements of the programs.

An aspect of the invention provides methods and systems for intelligent, destination task defined prioritization of inter-task communications (ITC) for a computer program, for architectural ITC performance isolation among a set of programs executing concurrently on a dynamically shared data processing platform, as well as for prioritizing instances of the program tasks for execution at least in part based on which of the instances have available to them their input data, including ITC data, enabling any given one of such instances to execute at the given time.

An aspect of the invention provides a system for prioritizing instances of a software program for execution. Such a system comprises: 1) a subsystem for determining which of the instances are ready to execute on an array of processing cores, at least in part based on whether a given one of the instances has available to it input data to process, and 2) a subsystem for assigning a subset of the instances for execution on the array of cores based at least in part on the determining. Various embodiments of that system include further features such as features whereby a) the input data is from a data source such that the given instance has assigned a high priority for purposes of receiving data; b) the input data is such data that it enables the given program instance to execute; c) the subset includes cases of none, some as well as all of the instances of said program; d) the instance is: a process, a job, a task, a thread, a method, a function, a procedure or an instance any of the foregoing, or an independent copy of the given program; and/or e) the system is implemented by hardware logic that is able to operate without software involvement.

An aspect of the invention provides a hardware logic implemented method for prioritizing instances of a software program for execution, with such a method involving: classifying instances of the program into the following classes, listed in the order from higher to lower priority for execution, i.e., in their reducing execution priority order: (I) instances indicated as having high priority input data for processing, and (II) any other instances. Various embodiments of that method include further steps and features such as features whereby a) the other instances are further classified into the following sub-classes, listed in their reducing execution priority order: (i) instances indicated as able to execute presently without the high priority input data, and (ii) any remaining instances; b) the high priority input data is data that is from a source where its destination instance,

of said program, is expecting high priority input data; c) a given instance of the program comprises tasks, with one of said tasks referred to as a destination task and others as source tasks of the given instance, and for the given instance, a unit of the input data is considered high priority if it is from such one of the source tasks that the destination task has assigned a high priority for inter-task communications to it; d) for any given one of the instances, a step of computing a number of its non-empty source task specific buffers among its input data buffers such that belong to source tasks of the given instance indicated at the time as high priority source tasks for communications to the destination task of the given instance, with this number referred to as an H number for its instance, and wherein, within the class I), the instances are prioritized for execution at least in part according to magnitudes of their H numbers, in descending order such that an instance with a greater H number is prioritized before an instance with lower H number; e) in case of two or more of the instances tied for the greatest H number, such tied instances are prioritized at least in part according to their respective total numbers of non-empty input data buffers, and/or f) at least one of the instances is either a process, a job, a task, a thread, a method, a function, a procedure, or an instance any of the foregoing, or an independent copy of the given program.

An aspect of the invention provides a system for processing a set of computer programs instances, with inter-task communications (ITC) performance isolation among the set of program instances. Such a system comprises: 1) a number of processing stages; and 2) a group of multiplexers connecting ITC data to a given stage among the processing stages, wherein a multiplexer among said group is specific to one given program instance among said set. The system hosts each task of the given program instance at different one of the processing stages, and supports copies of same task software code being located at more than one of the processing stages in parallel. Various embodiments of this system include further features such as a) a feature whereby at least one of processing stages comprises multiple processing cores such as CPU execution units, with, for any of the cores, at any given time, one of the program instances assigned for execution; b) a set of source task specific buffers for buffering data destined for a task of the given program instance located at the given stage, referred to as a destination task, and hardware logic for forming a hardware signal indicating whether sending ITC is presently permitted to a given buffer among the source task specific buffers, with such forming based at least in part on a fill level of the given buffer, and with such a signal being connected to a source task for which the given buffer is specific to; c) a feature providing, for the destination task, a set of source task specific buffers, wherein a given buffer is specific to one of the other tasks of the program instance for buffering ITC from said other task to the destination task; d) feature wherein the destination task provides ITC prioritization information for other tasks of the program instance located at their respective ones of the stages; d) a feature whereby the ITC prioritization information is provided by the destination task via a set of one or more hardware registers, with each register of the set specific to one of the other tasks of the program instance, and with each register configured to store a value specifying a prioritization level of the task that it is specific to, for purposes of ITC communications to the destination task; e) an arbitrator controlling from which source task of the program instance the multiplexer specific to that program instance will read its next ITC data unit for the destination task; and/or f) a feature whereby the arbi-

trator prioritizes source tasks of the program instance for selection by the multiplexer to read its next ITC data unit based at least in part on at least one of: (i) source task specific ITC prioritization information provided by the destination task, and (ii) source task specific availability information of ITC data for the destination task from the other tasks of the program instance.

Accordingly, aspects of the invention involve application-program instance specific hardware logic resources for secure and reliable ITC among tasks of application program instances hosted at processing stages of a multi-stage parallel processing system. Rather than seeking to inter-connect the individual processing stages or cores of the multi-stage manycore processing system as such, the invented mechanisms efficiently inter-connect the tasks of any given application program instance using the per application program instance specific inter-processing stage ITC hardware logic resources. Due to the ITC being handled with such application program instance specific hardware logic resources, the ITC performance experience by one application instance does not depend on the ITC resource usage (e.g. data volume and inter-task communications intensiveness) of the other applications sharing the given data processing system per the invention. This results in effective inter-application isolation for ITC in a multi-stage parallel processing system shared dynamically among multiple application programs.

An aspect of the invention provides systems and methods for scheduling instances of software programs for execution based at least in part on (1) availability of input data of differing priorities for any given one of the instances and/or (2) availability, on their fast-access memories, of memory contents needed by any given one of the instances to execute.

An aspect of the invention provides systems and methods for optimally allocating and assigning input port capacity to a data processing systems among data streams of multiple software programs based at least in part on input data load levels and contractual capacity entitlements of the programs.

An aspect of the invention provides systems and methods for resolution of resource access contentions, for resources including computing, storage and communication resources such as memories, queues, ports or processors. Such methods enable multiple potential user systems for a shared resource, in a coordinated and fair manner, to avoid conflicting resource access decisions, even while multiple user systems are deciding on access to set of shared resources concurrently, including at the same clock cycle.

An aspect of the invention provides systems and methods for load balancing, whereby the load balancer is configured to forward, by its first layer, any packets without destination instance within its destination application specified (referred to as no-instance-specified packets or NIS packets for short) it receives from its network input to such one of the processing systems in the local load balancing group that presently has the highest score for accepting NIS packets for the destination app of the given NIS packet. The load balancers further have destination processing system (i.e. for each given application, instance group) specific sub-modules, which, for NIS packets forwarded to them by the first layer balancing logic, specify a destination instance among the available, presently inactive instance resources of the destination app of a given NIS packet to which to forward the given NIS packet. In at least some embodiments of the invention, the score for accepting NIS packets for a destination processing system among the load balancing group is based at least in part on the amount of presently inactive

5

instance resources at the given processing system for the destination application of a given NIS packet.

#### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows, in accordance with an embodiment of the invention, a functional block diagram for a load balancing architecture for a bank of processor systems, such as those discussed in the following with reference to the remaining FIGS.

FIG. 2 shows, in accordance with an embodiment of the invention, a functional block diagram for a multi-stage manycore processing system shared dynamically among a set of software program instances, with the system providing capabilities for optimally scheduling inter-task communications (ITC) units between various tasks of any one of the program instances, as well as scheduling and placing instances of a given program task for execution on the processing stages of the system, at least in part based on which of the instances have available for them the input data, e.g. ITC data, needed by them to execute.

FIG. 3 shows, in accordance with an embodiment of the invention, a functional block diagram for a receive (RX) logic module of any of the processing stages of the multi-stage manycore processor system per FIG. 2.

FIG. 4 shows, in accordance with an embodiment of the invention, a functional block diagram for an application program specific submodule of the processing stage RX logic module per FIG. 3.

FIG. 5 shows, in accordance with an embodiment of the invention, a functional block diagram for an application program instance specific submodule of the application program specific submodule per FIG. 4.

FIG. 6 shows, in accordance with an embodiment of the invention, a functional block diagram for logic resources within one of the processing stages of a system 1 per FIG. 2 for connecting ITC data from input buffers of the RX logic (per FIGS. 3-5) to the manycore processor of the local processing stage.

FIG. 7 shows, in accordance with an embodiment of the invention, a functional block diagram for the application load adaptive manycore processor of a processing stage of the multi-stage processing system per preceding FIGS.

#### DETAILED DESCRIPTION

FIGS. and related descriptions in the following provide specifications for embodiments and aspects of hardware-logic based systems and methods for inter-task communications (ITC) with destination task defined source task prioritization, for input data availability based prioritization of instances of a given application task for execution on processing cores of a processing stage hosting the given task, for architecture-based application performance isolation for ITC in multi-stage manycore data processing system, as well as for load balancing of incoming processing data units among a group of such processing systems.

The invention is described herein in further detail by illustrating the novel concepts in reference to the drawings. General symbols and notations used in the drawings:

Boxes indicate a functional module comprising digital hardware logic.

Arrows indicate a digital signal flow. A signal flow may comprise one or more parallel bit wires. The direction of an arrow indicates the direction of primary flow of information associated with it with regards to discus-

6

sion of the system functionality herein, but does not preclude information flow also in the opposite direction.

A dotted line marks a border of a group of drawn elements that form a logical entity with internal hierarchy.

An arrow reaching to a border of a hierarchical module indicate connectivity of the associated information to/from all sub-modules of the hierarchical module.

Lines or arrows crossing in the drawings are decoupled unless otherwise marked.

For clarity of the drawings, generally present signals for typical digital logic operation, such as clock signals, or enable, address and data bit components of write or read access buses, are not shown in the drawings.

General notes regarding this specification (incl. text in the drawings):

For brevity: 'application (program)' is occasionally written in as 'app', 'instance' as 'inst' and 'application-task/instance' as 'app-task/inst' and so forth.

Terms software program, application program, application and program are used interchangeably in this specification, and each generally refers to any type of executable computer program.

In FIG. 5, and through the related discussions, the buffers 260 are considered to be First-in First-Out buffers (FIFO); however also other types than first-in first-out buffers can be used in various embodiments.

Illustrative embodiments and aspects of the invention are described in the following with references to the FIGS.

FIG. 1 presents the load balancing architecture for a row of processing systems per this description, comprising a set 4 of T load balancers 3 and a load balancing group 2 of S processing systems 1 (T and S are positive integers). Per this architecture, each of the balancers forward any no-instance-specific (NIS) packets (i.e. packets without a specific instance of their destination applications identified) arriving to them via their network inputs to one of the processing systems of the group, based on the NIS packet forwarding preference scores (for the destination app of the given NIS packet) of the individual processing systems of the load balancing group 2.

The load balancing per FIG. 1 for a bank 2 of the processing systems operates as follows:

The processing systems 1 count, for each of the application programs (apps) hosted on them:

a number X of their presently inactive instance resources, i.e., the number of additional parallel instances of the given app at the given processing system that could be activated at the time; and

from the above number, the portion Y (if any) of the additional activatable instances within the Core Entitlement (CE) level of the given app, wherein the CE is a number of processing cores at (any one of) the processing stages of the given processing system up to which the app in question is assured to get its requests for processing cores (to be assigned for its active instances) met;

the difference  $W=X-Y$ . The quantities X and/or W and Y, per each of the apps hosted on the load balancing group 2, are signaled 5 from each processing system 1 to the load balancers 4.

In addition, load balancing logic 4 computes the collective sum Z of the Y numbers across all the apps (with this across-apps-sum Z naturally being the same for all apps on a given processing system).

From the above numbers, for each app, the load balancer module 4 counts a no-instance-specified (NIS) packet

7

forwarding preference score (NIS score) for each processing system in the given load balancing group with a formula of:  $A*Y+B*W+C*Z$ , where A, B and C are software programmable, defaulting to e.g. A=4, B=1 and C=2.

In forming the NIS scores for a given app (by formula per above), a given instance of the app under study is deemed available for NIS packets at times that the app instance software has set an associated device register bit (specific to that app-inst) to an active value, and unavailable otherwise. The multiplexing (muxing) mechanism used to connect the app-instance software, from whichever core at its host manycore processor it may be executing at any given time, to its app-instance specific memory, is used also for connecting the app-instance software to its NIS-availability control device register.

The app-instance NIS availability control register of a given app-instance is reset (when the app-instance software otherwise would still keep its NIS availability control register at its active stage) also automatically by processing stage RX logic hardware whenever there is data at the input buffer for the given app-instance.

Each of the processing systems in the given load balancing group signals their NIS scores for each app hosted on the load balancing group to each of the load balancers **4** in front of the row **2** of processing systems. Also, the processing systems **1** provide to the load balancers app specific vectors (as part of info flows **9**) indicating which of their local instance resources of the given app are available for receiving NIS packets (i.e. packets with no destination instance specified).

Data packets from the network inputs **10** to the load balancing group include bits indicating whether any given packet is a NIS packet such that has its destination app but not any particular instance of the app specified. The load balancer **3** forwards any NIS packet it receives from its network input **10** to the processing system **1** in the local load balancing group **2** with the highest NIS score for the destination app of the given NIS packet. (In case of ties among the processing systems for the NIS score for the given destination app, the logic forwards the packet to the processing system among such tied systems based on their ID#, e.g. to the system with lowest ID#.) The forwarding of a NIS packet to a particular processing system **1** (in the load balancing group **2** of such systems) is done by this first layer of load balancing logic by forming packet write enable vectors where each given bit is a packet write enable bit specific to the processing system within the given load balancing group of the same system index # as the given bit in its write enable bit vector. For example, the processing system ID#**2** from a load balancing group of processing systems of ID#**0** through ID#**4** takes the bit at index **2** of the packet write enable vectors from the load balancers of the given group. In a straightforward scheme, the processing system #K within a given load balancing group hosts the instance group #K of each of the apps hosted by this group of the processing systems (where K=0, 1, . . . , max nr of processing systems in the load balancing group less 1).

The load balancers **3** further have destination processing system **1** (i.e. for each given app, instance group) specific submodules, which, for NIS packets forwarded to them by the first layer balancing logic (per above), specify a destination instance among the available

8

(presently inactive) instance resources of the destination app of a given NIS packet to which to forward the given NIS packet. In an straightforward scheme, for each given NIS packet forwarded to it, this instance group specific load balancing submodule selects, from the at-the-time available instances of the of the destination app, within the instance group that the given submodule is specific to, the instance resource with lowest ID#.

For other (not NIS) packets, the load balancer logic **3** simply forwards a given (non NIS) packet to the processing system **1** in the load balancing group **2** that hosts, for the destination app of the given packet, the instance group of the identified destination instance of the packet.

According to the forwarding decision per above bullet points, the (conceptual, actually distributed per the destination processing systems) packet switch module **6** filters packets from the output buses **15** of the load balancers **3** to input buses **19** of the destination processing systems, so that each given processing system **1** in the load balancing group **2** receives as active packet transmissions (marked e.g. by write by write enable signaling) on its input bus **19**, from the packets arriving from the load balancer inputs **10**, those packets that were indicated as destined to the given system **1** at entry to the load balancers, as well as the NIS packets that the load balancers of the set **4** forwarded to that given system **1**.

Note also that the network inputs **10** to the load balancers, as well as all the bold data path arrows in the FIGS., may comprise a number of parallel of (e.g. 10 Gbps) ports.

The load balancing logic implements coordination among port modules of the same balancer, so that any given NIS packet is forwarded, according to the above destination instance selection logic, to one of such app-instances that is not, at the time of the forwarding decision, already being forwarded a packet (incl. forwarding decisions made at the same clock cycle) by port modules with higher preference rank (e.g. based on lower port #) of the same balancer. Note that each processing system supports receiving packets destined for the same app-instance concurrently from different load balancers (as explained below).

The load balancers **3** support, per each app-inst, a dedicated input buffer per each of the external input ports (within the buses **10**) to the load balancing group. The system thus supports multiple packets being received (both via the same load balancer module **3**, as well as across the different load balancer modules per FIG. **1**) simultaneously for the same app-instances via multiple external input ports. From the load balancer input buffers, data packets are muxed to the processing systems **1** of the load balancing group so that the entry stage processor of each of the multi-stage systems (see FIG. **2**) in such group receives data from the load balancers similarly as the non-entry-stage processors receive data from the other processing stages of the given multi-stage processing system—i.e., in a manner that the entry stage (like the other stages) will get data per each of its app-instances at most via one of its input ports per a (virtual) source stage at any given time; the load balancer modules of the given load balancing group (FIG. **1**) appear thus as virtual source processing stages to entry stage of the multi-stage processing

systems of such load balancing group. The aforesaid functionality is achieved by logic at module 4 as detailed below:

To eliminate packet drops in cases where packets directed to same app-inst arrive in a time-overlapping manner through multiple input ports (within the buses 10) of same balancer 3, destination processing system 1 specific submodules at modules 3 buffer input data 15 destined for the given processing system 1 at app-inst specific buffers, and assign the processing system 1 input ports (within the bus 19 connecting to their associated processing system 1) among the app-insts so that each app-inst is assigned at any given time at most one input port per a load balancer 3. (Note that inputs to a processing system 1 from different load balancers 3 are handled by the entry stage (FIG. 2) the same way as the other processing stages 300 handle inputs from different source stages, as detailed in connection to FIG. 5—in a manner that supports concurrent reception of packets to the same destination app-inst from multiple source stages.) More specifically, the port capacity 19 for transfer of data from load balancers 4 to the given processing system 1 entry-stage buffers gets assigned using the same algorithm as is used for assignment of processing cores between the app-instances at the processing stages (FIG. 7), i.e., in a realtime input data load adaptive manner, while honoring the contractual capacity entitlements and fairness among the apps for actually materialized demands. This algorithm, which allocates at most one of the cores per each of the app-insts for the core allocation periods following each of its runs—and similarly assigns at most one of the ports at buses 19 to the given processing system 1 per each of the app-inst specific buffers queuing data destined for that processing system from any given source load balancer 3—is specified in detail in [1], Appendix A, Ch. 5.2.3. By this logic, the entry stage of the processing system (FIG. 2) will get its input data same way as the other stages, and there thus is no need to prepare for cases of multiple packets to same app-inst arriving simultaneously at any destination processing stage from any of its source stages or load balancers. This logic also ensures that any app with moderate input bandwidth consumption will get its contractually entitled share of the processing system input bandwidth (i.e. the logic protects moderate bandwidth apps from more input data intensive neighbors).

Note that since packet transfer within a load balancing group (incl. within the sub-modules of the processing systems) is between app-instance specific buffers, with all the overhead bits (incl. destination app-instance ID) transferred and buffered as parallel wires besides the data, core allocation period (CAP) boundaries will not break the packets while being transferred from the load balancer buffers to a given processing system 1 or between the processing stages of a given multi-stage system 1.

The mechanisms per above three bullet points are designed to eliminate all packet drops in the system such that are avoidable by system design, i.e., for reasons other than app-instance specific buffer overflows caused by systemic mismatches between input data loads to a given app-inst and the capacity entitlement level subscribed to by the given app.

FIG. 2 provides, according to an embodiment of the invention, a functional block diagram for a multistage many-core processor system 1 shared dynamically multiple concurrent application programs (apps), with hardware logic implemented capabilities for scheduling tasks of application program instances and prioritizing inter-task communications (ITC) among tasks of a given app instance, based at least in part on, for any given app-inst, at a given time, which tasks are expecting input data from which other tasks and which tasks are ready to execute on cores of the multi-stage manycore processing system, with the ready-to-execute status of a given task being determined at least in part based on whether the given task has available to it the input data from other tasks or system 1 inputs 19 so as to enable it to execute at the given time, including producing its processing outputs, such as ITC communications 20 to other tasks or program processing results etc. communications for external parties via external outputs 50. Operation and internal structure and elements of FIG. 2, other than for the aspects described herein, is, according to at least some embodiments of the invention, per [1], which the reader may review before this specification for context and background material.

In the architecture per FIG. 2, the multi-stage manycore processor system 1 is shared dynamically among tasks of multiple application programs (apps) and instances (insts) thereof, with, for each of the apps, each task located at one of the (manycore processor) based processing stages 300. Note however that, for any given app-inst, copies of same task software (i.e. copies of same software code) can be located at more than one of the processing stages 300 of a given system 1; thus the architecture per FIG. 2, with its any-to-any ITC connectivity between the stages 300, supports organizing tasks of a program flexibly for any desirable mixes or matches of pipelined and/or parallelized processing.

General operation of the application load adaptive, multi-stage parallel data processing system per FIG. 2, focusing on the main inputs to outputs data flows, is as follows: The system provides data processing services to be used by external parties (e.g. by clients of the programs hosted on the system) over networks. The system 1 receives data units (e.g. messages, requests, data packets or streams to be processed) from its users through its inputs 19, and transmits the processing results to the relevant parties through its network outputs 50. Naturally the network ports of the system of FIG. 2 can be used also for connecting with other (intermediate) resources and services (e.g. storage, databases etc.) as desired for the system to produce the requested processing results to the relevant external parties.

The application program tasks executing on the entry stage manycore processor are typically of 'master' type for parallelized/pipelined applications, i.e., they manage and distribute the processing workloads for 'worker' type tasks running (in pipelined and/or parallel manner) on the worker stage manycore processing systems (note that the processor system hardware is similar across all instances of the processing stages 300). The instances of master tasks typically do preliminary processing (e.g. message/request classification, data organization) and workflow management based on given input data units (packets), and then typically involve appropriate worker tasks at their worker stage processors to perform the data processing called for by the given input packet, potentially in the context of and in connection with other related input packets and/or other data elements (e.g. in memory or storage resources accessible by the system) referred to by such packets. (The processors have access to system memories through interfaces also additional to the IO

ports shown in FIG. 2, e.g. as described in [1], Appendix A, Ch. 5.4). Accordingly, the master tasks typically pass on the received data units (using direct connection techniques to allow most of the data volumes being transferred to bypass the actual processor cores) through the (conceptual) inter-stage packet-switch (PS) to the worker stage processors, with the destination application-task instance (and thereby, the destination worker stage) identified for each data unit as described in the following.

To provide isolation among the different applications configured to run on the processors of the system, by default the hardware controller of each processor 300, rather than any application software (executing on a given processor), inserts the application ID# bits for the data packets passed to the PS 200. That way, the tasks of any given application running on the processing stages in a system can trust that the packets they receive from the PS are from its own application. Note that the controller determines, and therefore knows, the application ID# that each given core within its processor is assigned to at any given time, via the application-instance to core mapping info that the controller produces. Therefore the controller is able to insert the presently-assigned app ID# bits for the inter-task data units being sent from the cores of its processing stage over the core-specific output ports to the PS.

While the processing of any given application (server program) at a system per FIG. 2 is normally parallelized and/or pipelined, and involves multiple tasks (many of which tasks and instances thereof can execute concurrently on the manycore arrays of the processing stages 300), the system enables external parties to communicate with any such application hosted on the system without knowledge about any specifics (incl. existence, status, location) of their internal tasks or instances. As such, the incoming data units to the system are expected to identify just their destination application, and when applicable, the application instance. Moreover, the system enables external parties to communicate with any given application hosted on a system through any of the network input ports 10 of any of the load balancers 3, without such external parties knowing whether or at which cores 520 (FIG. 7) or processing stages 300 any instance of the given application task (app-task) may be executing at any time.

Notably, the architecture enables the aforesaid flexibility and efficiency through its hardware logic functionality, so that no system or application software running on the system needs to either keep track of whether or where any of the instances of any of the app-tasks may be executing at any given time, or which port any given inter-task or external communication may have used. Thus the system, while providing a highly dynamic, application workload adaptive usage of the system processing and communications resources, allows the software running on and/or remotely using the system to be designed with a straightforward, abstracted view of the system: the software (both remote and local programs) can assume that all the applications, and all their tasks and instances, hosted on the given system are always executing on their virtual dedicated processor cores within the system. Also, where useful, said virtual dedicated processors can also be considered by software to be time-share slices on a single (unrealistically high speed) processor.

The presented architecture thereby enables achieving, at the same time, both the vital application software development productivity (simple, virtual static view of the actually highly dynamic processing hardware) together with high program runtime performance (scalable concurrent program

execution with minimized overhead) and resource efficiency (adaptively optimized resource allocation) benefits. Techniques enabling such benefits of the architecture are described in the following through more detailed technical description of the system 1 and its subsystems.

The any-to-any connectivity among the app-tasks of all the processing stages 300 provided by the PS 200 enables organizing the worker tasks (located at the array of worker stage processors) flexibly to suit the individual demands (e.g. task inter-dependencies) of any given application program on the system: the worker tasks can be arranged to conduct the work flow for the given application using any desired combinations of parallel and pipelined processing. E.g., it is possible to have the same task of a given application located on any number of the worker stages in the architecture per FIG. 2, to provide a desired number of parallel copies of a given task per an individual application instance, i.e. to support also data-parallelism, along with task concurrency.

The set of applications configured to run on the system can have their tasks identified by (intra-app) IDs according to their descending order of relative (time-averaged) workload levels. Under such (intra-app) task ID assignment principle, the sum of the intra-application task IDs, each representing the workload ranking of its tasks within its application, of the app-tasks hosted at any given processing system is equalized by appropriately configuring the tasks of differing ID#s, i.e. of differing workload levels, across the applications for each processing system, to achieve optimal overall load balancing. For instance, in case of T=4 worker stages, if the system is shared among M=4 applications and each of that set of applications has four worker tasks, for each application of that set, the busiest task (i.e. the worker task most often called for or otherwise causing the heaviest processing load among tasks of the app) is given task ID#0, the second busiest task ID#1, the third busiest ID#2, and the fourth ID#3. To balance the processing loads across the applications among the worker stages of the system, the worker stage #t gets task ID #t+m (rolling over at 3 to 0) of the application ID #m (t=0, 1, . . . T-1; m=0, 1, . . . M-1) (note that the master task ID#4 of each app is located at the entry/exit stages). In this example scenario of four application streams, four worker tasks per app as well as four worker stages, the above scheme causes the task IDs of the set of apps to be placed at the processing stages per Table 1 below:

TABLE 1

App ID # m (to right) Processing worker stage # t (below)	0	1	2	3
0	0	1	2	3
1	1	2	3	0
2	2	3	0	1
3	3	0	1	2

As seen in the example of Table 1, the sum of the task ID#s (with each task ID# representing the workload ranking of its task within its app) is the same for any row i.e. for each worker stage. This load balancing scheme can be straightforwardly applied for differing numbers of processing stages/tasks and applications, so that the overall task processing load is to be, as much as possible, equal across all worker-stage processors of the system. Advantages of such schemes include achieving optimal utilization efficiency of the processing resources and eliminating or at least mini-

13

mizing the possibility and effects of any of the worker-stage processors forming system-wide performance bottlenecks.

A non-exclusive alternative task to stage placement principle targets grouping tasks from the apps in order to minimize any variety among the processing core types demanded by the set of app-tasks placed on any given individual processing stage; that way, if all app-tasks placed on a given processing stage optimally run on the same processing core type, there is no need for reconfiguring the core slots of the manycore array at the given stage regardless which of the locally hosted app-tasks get assigned to which of its core slots (see [1], Appendix A, Ch. 5.5 for task type adaptive core slot reconfiguration, which may be used when the app-task located on the given processing stage demand different execution core types).

FIGS. 3-5 present the processing stage, app, app-instance level microarchitectures for the processing stage receive (RX) logic modules **201** (which collectively accomplish the functionality of the conceptual inter-stage packet-switch (PS) module of FIG. 2).

For a system of FIG. 2, note that the functionality of the conceptual inter-stage PS **200** is actually realized by instantiating the logic per FIG. 3 (and its submodules) as the RX logic of each manycore processing system **300** (referred to as a stage) in the multi-stage architecture; there is no need for other logic to the PS. Accordingly, in the hardware implementation, the stage RX logic **201** per FIG. 3-5 is part of the processing stage **300** that it interfaces to; i.e., in an actual hardware implementation, there is no PS module as its functionality is distributed to the individual processing stages.

Besides the division of the app-specific submodules **202** of the stage RX logic per FIG. 3 further to the array **410** of app-instance specific sub-modules **203**, FIG. 4 shows how the app-specific RX logic forms, for purposes of optimally assigning the processing cores of the local manycore processor among insts of the apps sharing the system, the following info for the given app:

Formation of a request for a number of processing cores (Core Demand Figure, CDF) at the local processing stage by the given app. The logic forms the CDF for the app based on the number of instances of the app that presently have (1) input data at their input buffers (with those buffers located at the instance specific stage RX logic submodules **203** per FIG. 5) and (2) their on-chip fast-access memory contents ready for the given instance to execute without access to the slower-access off-chip memories. In FIG. 4, (1) and (2) per above are signaled to the app-specific RX logic module **209** via the info flows **429** and **499** from the app-inst specific modules **203** (FIG. 5) and **800** (FIG. 7), respectively, per each of the insts of the app under study.

The priority order of instances of the app for purposes of selecting such instances for execution on the cores of the local manycore processor.

The info per the above two bullet points are sent from the RX logic **202** of each app via the info flow **430** to the controller **540** (FIG. 7) of the local manycore processor **500**, for the controller to assign optimal sets of the app-insts for execution on the cores **520** of the processor **500**.

The app-instance specific RX logic per FIG. 5 performs multiplexing **280** ITC packets from the source stage i.e. source task (of a given app-inst) specific First-in First-Out buffers (FIFOs) **260** to the local manycore processor via the input port **290** of that processor dedicated to the given app instance.

14

Note that when considering the case of RX logic of the entry-stage processing system of the multi-stage architecture per FIG. 4.1, note that in FIG. 5 and associated descriptions the notion of source stage/task naturally is replaced by the source load balancer, except in case of the ITC **20** from the exit stage to entry-stage, in which case the data source naturally is the exit stage processing system. However, the same actual hardware logic is instantiated for each occurrence of the processing stages **300** (incl. for the RX logic **201** of each stage) in this multi-stage architecture, and thus the operation of the stage RX logic can be fully explained by (as is done in the following) by assuming that the processing stage under study is instantiated as a worker or exit stage processing system, such that receives its input data from the other processing stages of the given multi-stage manycore processor, rather than from the load balancers of the given load balancing group, as in the case of the entry-stage processors; the load balancers appear to the entry-stage as virtual processing stages. Accordingly, when the RX logic of the entry stage manycore processor is considered, the references to 'source stage' are to be understood as actually referring to load balancers, and the references to ITC mean input data **19** to the multi-stage manycore processor system—except in case of the ITC **20** from the exit stage, as detailed above and as illustrated in FIG. 2. With this caveat, the description of the stage RX logic herein is written considering the operating context of worker and exit stage processors (with the same hardware logic being used also for the entry-stage).

Before the actual multiplexer, the app-instance specific RX logic per FIG. 5 has a FIFO module **245** per each of the source stages. The source-stage specific FIFO module comprises:

The actual FIFO **260** for queuing packets from its associated source stage that are destined to the local task of the app-instance that the given module per FIG. 5 is specific to.

A write-side multiplexer **250** (to the above referred FIFO) that (1) takes as its data inputs **20** the processing core specific data outputs **210** (see FIG. 7) from the processing stage that the given source-stage specific FIFO module is specific to, (2) monitors (via the data input overhead bits identifying the app-instance and destination task within it for any given packet transmission) from which one of its input ports **210** (within the bus **20**) it may at any given time be receiving a packet destined to the local task of the app-instance that the app-instance specific RX logic under study is specific to, with such an input referred to as the selected input, and (3) connects **255** to its FIFO queue **260** the packet transmission from the present selected input. Note that at any of the processing stages, at any given time, at most one processing core will be assigned for any given app instance. Thus any of the source stage specific FIFO modules **245** of the app-instance RX logic per FIG. 5 can, at any given time, receive data destined to the local task of the app-instance that the given app-instance RX logic module is specific to from at most one of the (processing core specific) data inputs of the write-side multiplexer (mux) **250** of the given FIFO module. Thus there is no need for separate FIFOs per each of the (e.g. 16 core specific) ports of the data inputs **20** at these source stage specific FIFO modules, and instead, just one common FIFO suffices per each given source stage specific buffering module **245**.

For clarity, the "local" task refers to the task of the app-instance that is located at the processing stage **300** that the

15

RX logic under study interfaces to, with that processing stage or processor being referred to as the local processing stage or processor. Please recall that per any given app, the individual tasks are located at separate processing stages. Note though that copies of the same task for a given app can be located at multiple processing stages in parallel. Note further that, at any of the processing stages, there can be multiple parallel instances of any given app executing concurrently, as well as that copies of the task can be located in parallel at multiple processing stages of the multi-stage architecture, allowing for processing speed via parallel execution at application as well as task levels, besides between the apps.

The app-instance RX module **203** per FIG. **5** further provides arbitrating logic **270** to decide, at multiplexing packet boundaries **281**, from which of the source stage FIFO modules **245** to mux **280** out the next packet to the local manycore processor via the processor data input port **290** specific to the app-instance under study. This muxing process operates as follows:

Each given app-instance software provides a logic vector **595** to the arbitrating logic **270** of its associated app-instance RX module **203** such that has a priority indicator bit within it per each of its individual source stage specific FIFO modules **245**: while a bit of such a vector relating to a particular source stage is at its active state (e.g. logic '1'), ITC from the source stage in question to the local task of the app-instance will be considered to be high priority, and otherwise normal priority, by the arbitrator logic in selecting the source stage specific FIFO from where to read the next ITC packet to the local (destination) task of the studied app-instance.

The arbitrator selects the source stage specific FIFO **260** (within the array **240** of the local app-instance RX module **203**) for reading **265**, **290** the next packet per the following source priority ranking algorithm:

The source priority ranking logic maintains three logic vectors as follows:

- 1) A bit vector wherein each given bit indicates whether a source stage of the same index as the given bit is both assigned by the local (ITC destination) task of the app-instance under study a high priority for ITC to it and has its FIFO **260** fill level above a configured monitoring threshold;
- 2) A bit vector wherein each given bit indicates whether a source stage of the same index as the given bit is both assigned a high priority for ITC (to the task of the studied app-instance located at the local processing stage) and has its FIFO non-empty;
- 3) A bit vector wherein each given bit indicates whether a source stage of the same index as the given bit has its FIFO fill level above the monitoring threshold; and
- 4) A bit vector wherein each given bit indicates whether a source stage of the same index as the given bit has data available for reading.

The FIFO **260** fill level and data-availability is signaled in FIG. **5** via info flow **261** per each of the source-stage specific FIFO modules **245** of the app-inst specific array **240** to the arbitrator **270** of the app-inst RX module, for the arbitrator, together with the its source stage prioritization control logic **285**, to select **272** the next packet to read from the optimal source-stage specific FIFO module **245** (as detailed below).

The arbitrator logic **270** also forms (by logic OR) an indicator bit for each of the above vectors 1) through 4) telling whether the vector associated with the given

16

indicator has any bits in its active state. From these indicators, the algorithm searches the first vector, starting from vector 1) and proceeding toward vector 4), that has one or more active bits; the logic keeps searching until such a vector is detected.

From the detected highest priority ranking vector with active bit(s), the algorithm scans bits, starting from the index of the current start-source-stage (and after reaching the max bit index of the vector, continuing from bit index **0**), until it finds a bit in an active state (logic '1'); the index of such found active bit is the index of the source stage from which the arbitrator controls its app-instance port mux **280** to read **265** its next ITC packet for the local task of the studied app-instance.

The arbitrator logic uses a revolving (incrementing by one at each run of the algorithm, and returning to 0 from the maximum index) starting source stage number as a starting stage in its search of the next source stage for reading an ITC packet.

When the arbitrator has the appropriate data source (from the array **240**) thus selected for reading **265**, **290** the next packet, the arbitrator **270** directs **272** the mux **280** to connect the appropriate source-stage specific signal **265** to its output **290**, and accordingly activates, when enabled by the read-enable control **590** from the app-inst software, the read enable **271** signal for the FIFO **260** of the presently selected source-stage specific module **245**.

Note that the ITC source task prioritization info **595** from the task software of app-instances to their RX logic modules **203** can change dynamically, as the processing state and demands of input data for a given app-instance task evolve over time, and the arbitrator modules **270** (FIG. **5**) apply the current state of the source task prioritization info provided to them in selecting from which of the source stages to multiplex **280** out the next ITC packet over the output port **290** of the app-instance RX logic. In an embodiment, the local task of a given app-inst, when a need arises, writes **575**, **595** the respective ITC prioritization levels for its source tasks (of the given app-inst) on its source-task specific ITC prioritization hardware registers, which are located at (or their info connected to) source-stage prioritization control logic submodule **285** of the arbitrator **270** of the RX module **203** of that given app-inst. Please see FIG. **7** for the muxing **580** of the input data read control info (incl. source prioritization) from the app-insts executing at the cores of the array to their associated RX modules **203**.

In addition, the app-instance RX logic per FIG. **5** participates in the inter-stage ITC flow-control operation as follows:

Each of the source stage specific FIFO modules **245** of a given app-instance at the RX logic for a given processing stage maintains a signal **212** indicating whether the task (of the app instance under study) located at the source stage that the given FIFO **260** is specific to is presently permitted to send ITC to the local (destination) task of the app-instance under study: the logic denies the permit when the FIFO fill level is above a defined threshold, while it otherwise grants the permit.

As a result, any given (source) task, when assigned for execution at a core **520** (FIG. **7**) at the processing stage where the given task is located, receives the ITC sending permission signals from each of the other (destination) tasks of its app-instance. Per FIG. **7**, these ITC permissions are connected **213** to the processing cores of the (ITC source) stages through multiplexers **600**, which, according to the control **560** from the controller **540** at the given (ITC source) processing stage identifying the active app-instance for each



execution core **520**, connect **213** the incoming ITC permission signals **212** from the other stages of the given multi-stage system **1** to the cores **520** at that stage. For this purpose, the processing stage provides core specific muxes **600**, each of which connects to its associated core the incoming ITC send permit signals from the 'remote' (destination) tasks of the app-instance assigned at the time to the given core, i.e., from the tasks of that app-instance located at the other stages of the given processing system. The (destination) task RX logic modules **203** activate the ITC permission signals for times that the source task for which the given permission signal is directed to is permitted to send further ITC data to that destination task of the given app-inst.

Each given processing stage receive and monitor ITC permit signal signals **212** from those of the processing stages that the given stage actually is able to send ITC data to; please see FIG. 2 for ITC connectivity among the processing stages in the herein studied embodiment of the presented architecture.

The ITC permit signal buses **212** will naturally be connected across the multi-stage system **1** between the app-instance specific modules **203** of the RX logic modules **202** of the ITC destination processing stages and the ITC source processing stages (noting that a given stage **300** will be both a source and destination for ITC as illustrated in FIG. 2), though the inter-stage connections of the ITC flow control signals are not shown in FIG. 2. The starting and ending points of the of the signals are shown, in FIG. 5 and FIG. 7 respectively, while the grouping of these ITC flow control signals according to which processing stage the given signal group is directed to, as well as forming of the stage specific signal groups according to the app-instance # that any given ITC flow control signal concerns, are illustrated also in FIGS. 3-4. In connecting these per app-instance ID# arranged, stage specific groups of signals (FIG. 3) to any of the processing stages **300** (FIG. 7), the principle is that, at arrival to the stage that a given set of such groups of signals is directed to, the signals from said groups are re-grouped to form, for each of the app-instances hosted on the system **1**, a bit vector where a bit of a given index indicates whether the task of a given app-instance (that the given bit vector is specific to) hosted at this (source) stage under study is permitted at that time to send ITC data to its task located at the stage ID# of that given index. Thus, each given bit in these bit vectors informs whether the studied task of the given app-instance is permitted to send ITC to the task of that app-instance with task ID# equal to the index of the given bit. With the incoming ITC flow control signals thus organized to app-instance specific bit vectors, the above discussed core specific muxes **600** (FIG. 7) are able to connect to any given core **520** of the local manycore array the (task-ID-indexed) ITC flow control bit vector of the app-instance presently assigned for execution at the given core. By monitoring the destination stage (i.e. destination task) specific bits of the ITC permission bit vector thus connected to the present execution core of a task of the studied app-instance located at the ITC (source) processing stage under study (at times that the given app-instance actually is assigned for execution), that ITC source task will be able to know to which of the other tasks of its app-instance sending ITC is permitted at any given time.

Note that, notwithstanding the functional illustration in FIG. 5, in actual hardware implementation, the FIFO fill-above-threshold indications from the source stage specific FIFOs **260** of the app-instance specific submodules of the RX logic modules of the (ITC destination) processing stages

of the present multi-stage system are wired directly, though as inverted, as the ITC send permission indication signals to the appropriate muxes **600** of the (ITC source) stages, without going through the arbitrator modules (of the app-instance RX logic modules at the ITC destination stages). Naturally, an ITC permission signal indicating that the destination FIFO for the given ITC flow has its fill level presently above the configured threshold is to be understood by the source task for that ITC flow as a denial of the ITC permission (until that signal would turn to indicate that the fill level of the destination FIFO is below the configured ITC permission activation threshold).

Each source task applies these ITC send permission signals from a given destination task of its app-instance at times that it is about to begin sending a new packet over its (assigned execution core specific) processing stage output port **210** to that given destination task. The ITC destination FIFO **260** monitoring threshold for allowing/disallowing further ITC data to be sent to the given destination task (from the source task that the given FIFO is specific to) is set to a level where the FIFO still has room for at least one ITC packet worth of data bytes, with the size of such ITC packets being configurable for a given system implementation, and the source tasks are to restrict the remaining length of their packet transmissions to destination tasks denying the ITC permissions according to such configured limits.

The app-level RX logic per FIG. 4 arranges the instances of its app for the instance execution priority list **535** (sent via info flow **430**) according to their descending order of their priority scores computed for each instance based on their numbers **429** of source stage specific non-empty FIFOs **260** (FIG. 5) as follows. To describe the forming of priority scores, we first define (a non-negative integer)  $H$  as the number of non-empty FIFOs of the given instance whose associated source stage was assigned a high ITC priority (by the local task of the given app-instance hosted at the processing stage under study). We also define (a non-negative integer)  $L$  as the number of other (non-high ITC priority source task) non-empty FIFOs of the given instance. With  $H$  and  $L$  thus defined, the intra-app execution priority score  $P$  for a given instance specific module (of the present app under study) is formed with equations as follows, with different embodiments having differing coefficients for the factors  $H$ ,  $L$  and the number of tasks for the app,  $T$ :

for  $H > 0$ ,  $P = T - 1 + 2H + L$ ; and  
for  $H = 0$ ,  $P = L$ .

The logic for prioritizing the instances of the given app for its execution priority list **535**, via a continually repeating process, signals (via hardware wires dedicated for the purpose) to the controller **540** of the local manycore processor **500** (FIG. 7) this instance execution priority list using the following format:

The process periodically starts from priority order 0 (i.e. the app's instance with the greatest priority score  $P$ ), and steps through the remaining priority orders 1 through the maximum supported number of instances for the given application (specifically, for its task located at the processing stage under study) less 1, producing one instance entry per each step on the list that is sent to the controller as such individual entries. Each entry of such a priority list comprises, as its core info, simply the instance ID# (as the priority order of any given instance is known from the number of clock cycles since the bit pulse marking the priority order 0 at the start of a new list). To simplify the logic, also the priority order (i.e. the number of clock cycles since the bit pulse marking the priority order 0) of any given entry on these lists is sent along with the instance ID#.

At the beginning of its core to app-instance assignment process, the controller **540** of the manycore processor uses the most recent set of complete priority order lists **535** received from the application RX modules **202** to determine which (highest priority) instances of each given app to assign for execution for the next core allocation period on that processor.

Per the foregoing, the ITC source prioritization, program instance execution prioritization and ITC flow control techniques provide effective program execution optimization capabilities for each of a set of individual programs configured to dynamically share a given data processing system **1** per this description, without any of the programs impacting or being impacted by in any manner the other programs of such set. Moreover, for ITC capabilities, also the individual instances (e.g. different user sessions) of a given program are fully independent from each other. The herein described techniques and architecture thus provide effective performance and runtime isolation between individual programs among groups of programs running on the dynamically shared parallel computing hardware.

From here, we continue by exploring the internal structure and operation of a given processing stage **300** beyond its RX logic per FIGS. **3-5**, with references to FIGS. **6** and **7**.

Per FIG. **6**, any of the processing stages **300** of the multi-stage system **1** per FIG. **2** has, besides the RX logic **201** and the actual manycore processor system (FIG. **7**), an input multiplexing subsystem **450**, which connects input data packets from any of the app-instance specific input ports **290** to any of the processing cores **520** of the processing stage, according to which app-instance is executing at any of the cores at any given time.

The monitoring of the buffered input data availability **261** at the destination app-instance FIFOs **260** of the processing stage RX logic enables optimizing the allocation of processing core capacity of the local manycore processor among the application tasks hosted on the given processing stage. Since the controller module **540** of the local manycore processor determines which instances of the locally hosted tasks of the apps in the system **1** execute at which of the cores of the local manycore array **515**, the controller is able to provide the dynamic control **560** for the muxes **450** per FIG. **6** to connect the appropriate app-instance specific input data port **290** from the stage RX logic to each of the core specific input data ports **490** of the manycore array of the local processor.

Internal elements and operation of the application load adaptive manycore processor system **500** are illustrated in FIG. **7**. For the intra processing stage discussion, it shall be recalled that there is no more than one task located per processing stage per each of the apps, though there can be up to X (a positive integer) parallel instances of any given app-task at its local processing stage (having an array **515** of X cores). With one task per application per processing stage **300**, the term app-instance in the context of a single processing stage means an instance of an app-task hosted at the given processing stage under study.

FIG. **7** provides a functional block diagram for the manycore processor system dynamically shared among instances of the locally hosted app-tasks, with capabilities for application input data load adaptive allocation of the cores **520** among the applications and for app-inst execution priority based assignment of the cores (per said allocation), as well as for accordantly dynamically reconfigured **550**, **560** I/O and memory access by the app-insts.

As illustrated in FIG. **7**, the processor system **500** comprises an array **515** of processing cores **520**, which are dynamically shared among instances of the locally hosted

tasks of the application programs configured to run on the system **1**, under the direction **550**, **560** of the hardware logic implemented controller **540**. Application program specific logic functions at the RX module (FIG. **3-5**) signal their associated applications' capacity demand indicators **430** to the controller. Among each of these indicators, the core-demand-figures (CDFs) **530**, express how many cores their associated app is presently able utilize for its (ready to execute) instances. Each application's capacity demand expressions **430** for the controller further include a list of its ready instances in an execution priority order **535**.

Any of the cores **520** of a processor per FIG. **7** can comprise any types of software program and data processing hardware resources, e.g. central processing units (CPUs), graphics processing units (GPUs), digital signal processors (DSPs) or application specific processors (ASPs) etc., and in programmable logic (FPGA) implementation, the core type for any core slot **520** is furthermore reconfigurable per expressed demands of its assigned app-task, e.g. per [1], Appendix A, Ch. 5.5.

The hardware logic based controller **540** module within the processor system, through a periodic process, allocates and assigns the cores **520** of the processor among the set of applications and their instances based on the applications' core demand figures (CDFs) **530** as well as their contractual core capacity entitlements (CEs). This application instance to core assignment process is exercised periodically, e.g. at intervals such as once per a defined number (for instance 64, 256 or 1024, or so forth) of processing core clock or instruction cycles. The app-instance to core assignment algorithms of the controller produce, per the app-instances on the processor, identification **550** of their execution cores (if any, at any given time), as well as per the cores of the fabric, identification **560** of their respective app-instances to execute. Moreover, the assignments **550**, **560** between app-insts and the cores of the array **515** control the access between the cores **520** of the fabric and the app-inst specific memories at the fabric network and memory subsystem **800** (which can be implemented e.g. per [1] Appendix A, Ch. 5.4).

The app-instance to core mapping info **560** also directs the muxing **450** of input data from the RX buffers **260** of an appropriate app-instance to each core of the array **515**, as well as the muxing **580** of the input data read control signals (**570** to **590**, and **575** to **595**) from the core array to the RX logic submodule (FIG. **5**) of the app-instance that is assigned for any given core **520** at any given time.

Similarly, the core to app-inst mapping info **560** also directs the muxing **600** of the (source) app-instance specific ITC permit signals (**212** to **213**) from the destination processing stages to the cores **520** of the local manycore array, according to which app-instance is presently mapped to which core.

Further reference specifications for aspects and embodiments of the invention are in the references [1] through [11].

The functionality of the invented systems and methods described in this specification, where not otherwise mentioned, is implemented by hardware logic of the system (wherein hardware logic naturally also includes any necessary signal wiring, memory elements and such).

Generally, this description and drawings are included to illustrate architecture and operation of practical embodiments of the invention, but are not meant to limit the scope of the invention. For instance, even though the description does specify certain system elements to certain practical types or values, persons of skill in the art will realize, in view of this description, that any design utilizing the architectural

21

or operational principles of the disclosed systems and methods, with any set of practical types and values for the system parameters, is within the scope of the invention. Moreover, the system elements and process steps, though shown as distinct to clarify the illustration and the description, can in various embodiments be merged or combined with other elements, or further subdivided and rearranged, etc., without departing from the spirit and scope of the invention. Finally, persons of skill in the art will realize that various embodiments of the invention can use different nomenclature and terminology to describe the system elements, process phases etc. technical concepts in their respective implementations. Generally, from this description many variants and modifications will be understood by one skilled in the art that are yet encompassed by the spirit and scope of the invention.

What is claimed is:

1. A system for prioritizing instances of a software program for execution, the system including:

a subsystem for determining which of the instances are ready to execute on an array of processing cores, at least in part based on whether a given one of the instances has available input data to process, wherein said input data: (a) is sent to the given instance of the software program (i) by an instance of a different task of the same program or (ii) by an instance of a different software program, and (b) is queued at data source specific buffers of the given instance, and the given instance allots input data receive priorities for the data sources associated with said buffers;

a subsystem for computing a priority score for each of the instances, wherein the priority score for a given one of the instances is computed at least in part based on a weighted sum of (1) a number of non-empty source-specific buffers of the given instance whose associated data source was allotted a high input data receive priority by the given instance, and (2) a number of other non-empty source-specific buffers of the given instance; and

a subsystem for assigning a subset of the instances for execution on the array of cores at least in part so that instances with highest priority scores get selected for the subset assigned for execution, wherein the subset includes cases of none, some as well as all of the instances of said program,

wherein said system is implemented in hardware logic.

2. The system of claim 1 wherein the input data is such data that enables the given program instance to execute.

3. The system of claim 1 implemented by hardware logic that is configured to operate without software involvement.

4. A method for prioritizing instances of a software program task for execution, with said task considered herein as a destination task for its source tasks, the method involving:

classifying the instances into the following classes, listed in their reducing execution priority order:

- 1) instances indicated as having high priority input data for processing, with said high priority assigned by such one of the instances that said data is to be processed by; and
- 2) other instances;

based at least in part on said classifying, executing at least some of the instances on a set of hardware processing resources; and

for any given one of the instances, computing a number of its non-empty source task specific buffers among input data buffers of the given instance such that belong to source tasks of the given instance indicated at the

22

time as high priority source tasks for communications to the given instance, with this number referred to as an H number for its instance, wherein, within the class 1), the instances are prioritized for execution at least in part according to magnitudes of their H numbers, in descending order such that an instance with a greater H number is prioritized for execution before an instance with lower H number.

5. The method of claim 4 wherein the other instances are further classified into the following sub-classes, listed in their reducing execution priority order:

- a) instances indicated as able to execute presently without the high priority input data; and
- b) other instances.

6. The method of claim 4 wherein a unit of the input data is considered high priority if the unit of the input data is from such one of the source tasks that the given instance has assigned a high priority for inter-task communications destined to itself.

7. The method of claim 4 wherein, in case of two or more of the instances tied for the greatest H number, such tied instances are prioritized at least in part according to their respective total numbers of non-empty input data buffers.

8. The method of claim 4 implemented by hardware logic.

9. The method of claim 4 wherein the task is chosen from: a process, a job, an actor, a thread, a method, a function, and a procedure.

10. A method for prioritizing instances of a software program for execution, the method including:

determining which of the instances are ready to execute on an array of processing cores, at least in part based on whether a given one of the instances has available input data to process, wherein said input data: (a) is sent to the given instance of the software program (i) by an instance of a different task of the same program or (ii) by an instance of a different software program, and (b) is queued at data source specific buffers of the given instance, and the given instance allots input data receive priorities for the data sources associated with said buffers;

computing a priority score for each of the instances, wherein the priority score for a given one of the instances is computed at least in part based on a weighted sum of (1) a number of non-empty source-specific buffers of the given instance whose associated data source was allotted a high input data receive priority by the given instance, and (2) a number of other non-empty source-specific buffers of the given instance; and

assigning a subset of the instances for execution on the array of cores at least in part so that instances with highest priority scores get selected for the subset assigned for execution, wherein the subset includes cases of none, some as well as all of the instances of said program,

wherein said method is implemented by hardware logic.

11. The method of claim 10 wherein the input data is such data that enables the given program instance to execute.

12. The method of claim 10 implemented by hardware logic that is configured to operate without software involvement.

13. A system for prioritizing instances of a software program task for execution, with said task considered herein as a destination task for its source tasks, the method involving:

**23**

a subsystem for classifying the instances into the following classes, listed in their reducing execution priority order:

- 1) instances indicated as having high priority input data for processing, with said high priority assigned by such one of the instances that said data is to be processed by; and

- 2) other instances;

a subsystem for executing, at least in part based on said classifying, at least some of the instances on a set of hardware processing resources; and

a subsystem for computing, for any given one of the instances, a number of its non-empty source task specific buffers among input data buffers of the given instance such that belong to source tasks of the given instance indicated at the time as high priority source tasks for communications to the given instance, with this number referred to as an H number for its instance, wherein, within the class 1), the instances are prioritized for execution at least in part according to magnitudes of their H numbers, in descending order such that an instance with a greater H number is prioritized for execution before an instance with lower H number.

**24**

**14.** The system of claim **13** wherein the other instances are further classified into the following sub-classes, listed in their reducing execution priority order:

- a) instances indicated as able to execute presently without the high priority input data; and
- b) other instances.

**15.** The system of claim **13** wherein a unit of the input data is considered high priority if the unit of the input data is from such one of the source tasks that the given instance has assigned a high priority for inter-task communications destined to itself.

**16.** The system of claim **13** wherein, in case of two or more of the instances tied for the greatest H number, such tied instances are prioritized at least in part according to their respective total numbers of non-empty input data buffers.

**17.** The system of claim **13** implemented in hardware logic.

**18.** The system of claim **13** wherein the task is chosen from: a process, a job, an actor, a thread, a method, a function, and a procedure.

\* \* \* \* \*